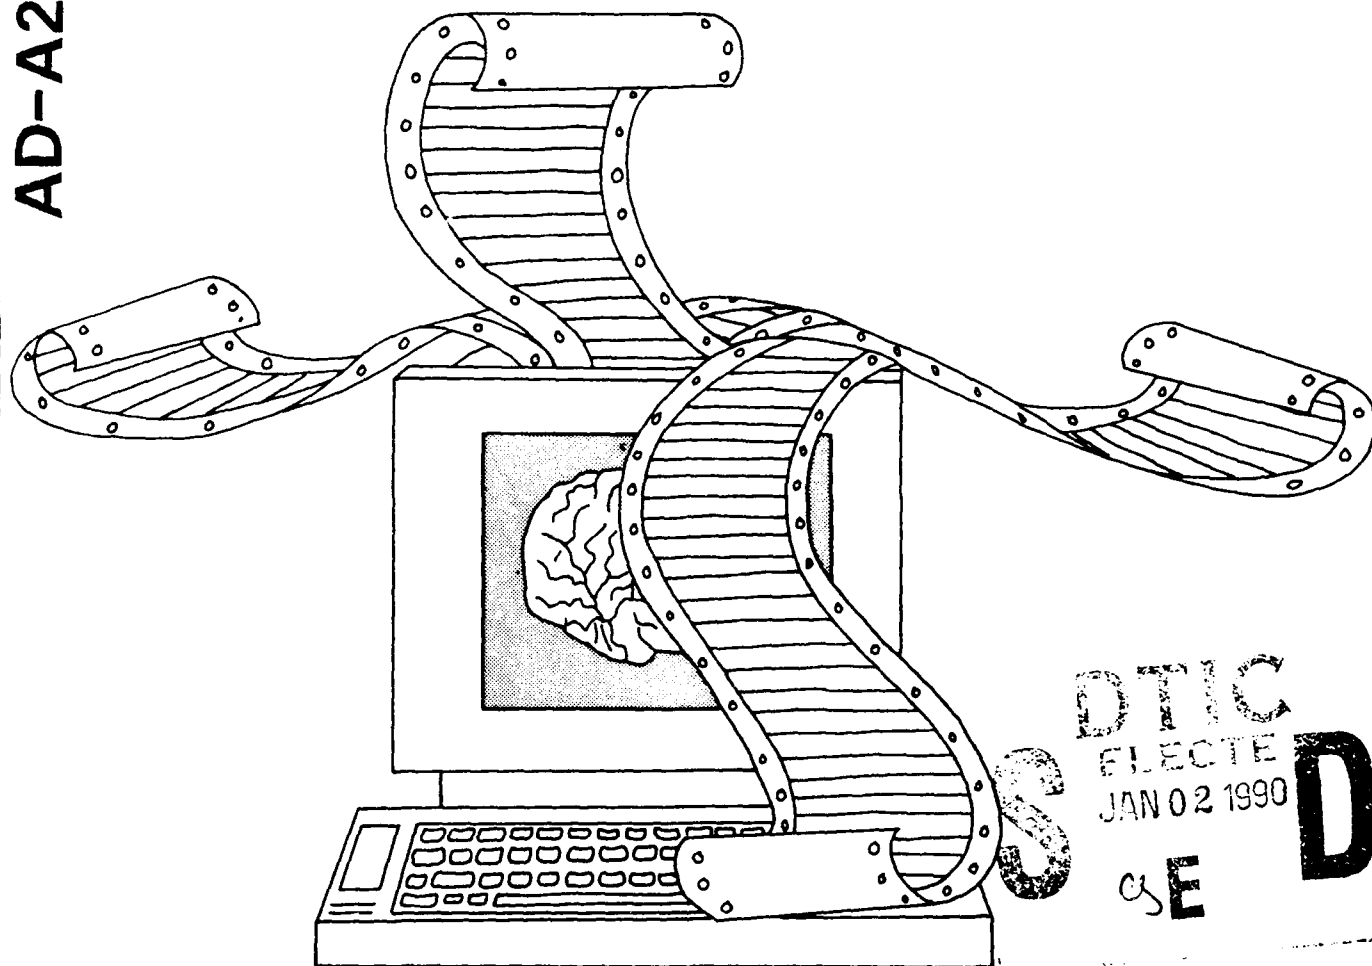


AD-A215 915

DTIC FILE COPY

2



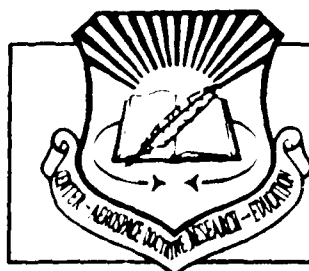
DTIC
ELECTE
JAN 02 1990
E

D

ARTIFICIAL INTELLIGENCE : A "USER FRIENDLY" INTRODUCTION

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited



AIR UNIVERSITY (AU)
MAXWELL AIR FORCE BASE, ALABAMA

90 01 02 004

After you have read the research report, please give us your frank opinion on the contents. All comments--large or small, complimentary or caustic--will be gratefully appreciated. Mail them to: CADRE/RI, Maxwell AFB AL 36112.



Cut Along Dotted Line

Thank you for your assistance



Research Report No. AU-ARI-85-1

ARTIFICIAL INTELLIGENCE
A "USER FRIENDLY"
INTRODUCTION

by

Pat O. Clifton
Colonel, USAF
Research Fellow, ARI
Center for Aerospace Doctrine, Research, and Education

Air University
Air University Press
Maxwell Air Force Base, Alabama 36112-5532

March 1985

Accession For	
DTIC	<input checked="checked" type="checkbox"/>
DDI	<input type="checkbox"/>
U. S. Armed	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

DISCLAIMER

This study represents the views of the author and does not necessarily reflect the official opinion of the Air University Center for Aerospace Doctrine, Research, and Education (CADRE) or the Department of the Air Force. This manuscript has been reviewed by security and policy review authorities and is cleared for public release.

This document is the property of the United States Government and is not to be reproduced in whole or in part without permission of the Commander, CADRE, Maxwell Air Force Base, Alabama.

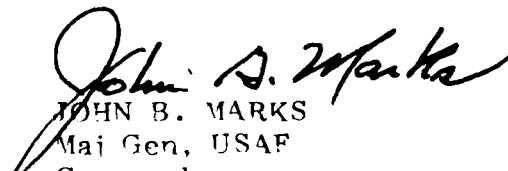
CONTENTS

<u>Chapter</u>		<u>Page</u>
	DISCLAIMER-----	ii
	FOREWORD-----	v
	ABOUT THE AUTHOR-----	vii
	PREFACE-----	ix
	INTRODUCTION-----	xi
1	BACKGROUND INFORMATION-----	1
	Introduction-----	1
	Definitions-----	2
	Origins of AI-----	4
	Conventional Computers-----	7
	Basic Computer Operations-----	8
	AI Computer Systems-----	9
	System Comparison-----	10
	AI Research Progress-----	12
	Notes-----	14
2	ARTIFICIAL INTELLIGENCE APPLICATION-----	17
	Expert Systems-----	17
	System Components-----	18
	Search Procedures-----	22
	Existing Expert Systems-----	23
	Natural Language-----	24
	Natural Language Interface System-----	25
	Translation Projects-----	26
	Voice Recognition-----	27
	Vision Systems-----	29
	Robotics-----	31
	Development Problems-----	32
	Notes-----	35
3	FUTURE PROSPECTS-----	37
	Overview-----	37
	Fifth Generation Projects-----	39
	Japan's Fifth Generation Computer Project-----	39
	US Fifth Generation Projects-----	41
	Government Projects-----	42
	Commercial AI Projects-----	44
	Great Britain-----	45
	Common Market Project-----	46
	Soviet Fifth Generation Project-----	46
	Potential Applications for AI Technology-----	47
	Military Applications-----	49
	Natural Language-----	50
	Civilian Applications-----	52

	<u>Page</u>
Summary and Conclusions-----	54
Notes-----	58
ATTACHMENT 1-----	61

FOREWORD

The Electronic Security Command, like other government agencies, is studying the potential applications of artificial intelligence (AI). Our initial look at AI research indicated that we should first pursue an in-depth review of the state of AI development. As part of that study we tasked our Command-sponsored Research Fellow at Air University for 1984-85, Col Pat O. Clifton, to prepare an introductory handbook on AI. We believe it essential that senior members of our staff become more familiar with AI's rapidly evolving technology. Artificial intelligence has become a popular "buzzword" among those working with computer support systems. It is a complex subject offering potential capabilities that are not clearly understood. What this study provides is an introduction to artificial intelligence for senior civilian and military leaders. It should help clear away some of the promotional "hype" that clouds what AI really can and can't do. This project was completed in coordination with Air University's new Center for Aerospace Doctrine, Research, and Education (CADRE). We at ESC believe CADRE will continue to provide critical support to our future operational needs.


JOHN B. MARKS
Maj Gen, USAF
Commander
Electronic Security Command

ABOUT THE AUTHOR

Col Pat O. Clifton was the Electronic Security Command's (ESC) Research Fellow at CADRE for AY 84-85. He is a graduate of Squadron Officers School and the Air War College, and a distinguished military graduate of the Air Command and Staff College. He received a BA degree from Oklahoma State University and an MA from Creighton University. Col Clifton is a career intelligence officer whose assignments include tours as an intelligence support squadron commander and a wing operations officer. Col Clifton was recently selected to be the commander of the 3480 Technical Training Group, Goodfellow AFB.

PREFACE

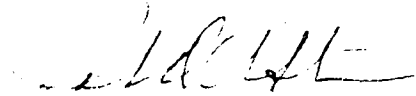
CADRE's value to the Air Force is the collection of academic skills it can focus on any given research project. This unique center brings together civilian and military research fellows who were extremely helpful in reviewing and providing comments on this handbook. The entire Air University Press (AUP) editorial staff also helped with guidance on research and writing procedures. I must single out Mr Preston Bryant and Dr Paul Godwin who served as my primary editors. Their suggestions always improved the clarity of the text while eliminating needless verbiage. I also want to thank the AUP document processing staff. Led by Marcia Williams, they worked miracles turning a mishmash of scribbles into a finished product.

Dr John Romo, who received his PhD in Mathematics from Oklahoma State University, provided invaluable technical advice and assistance on the entire paper. He is assigned to the Systems Technology Directorate at HQ ESC.

Dr Alan Roecks, who received his PhD from the University of Wisconsin-Madison, in Computers in Education, provided numerous source materials and editorial comments. He is also assigned to the Systems Technology Directorate at HQ ESC.

This handbook would never have been started, much less completed, without continuing support from Col Rolf Smith, Director of Plans and Requirements, Headquarters Electronic Security Command (ESC). Colonel Smith is a PhD candidate (ABD) in computer science at Texas A&M University. He suggested the research topic and ensured that it received headquarters support during its development. Colonel Smith

provided the kind of major command support that CADRE must have to succeed. Any ultimate success this handbook might enjoy is a testament to the CADRE concept. It is working.



PAT O. CLIFTON
Colonel, USAF
Research Fellow
Airpower Research Institute

INTRODUCTION

Artificial Intelligence (AI) may be one of the most promising and misunderstood technologies evolving today. The term artificial intelligence often connotes super-human computers, mysterious black boxes, and mechanical wizards. In fact, developments in AI are based on almost three decades of deliberate research. With the advent of the first electronic 'calculating machines' in the mid-1940s, AI pioneers realized that machines could be made to manipulate symbols and concepts as well as numerical values. The search for 'intelligent' systems was on. Early 'smart' computers were taught to play games such as checkers and chess. Research progressed to other types of problem solving systems. Early attempts to use AI techniques in language translations were such failures, however, that AI research in general suffered. In the late seventies, new computer technologies revived interest in AI. New very large scale integrated (VLSI) chips, for example, provided the memory capacity necessary for successful AI applications. Both government and commercial organizations began renewed efforts to exploit the new promise of AI. In 1981, the Japanese announced their intention of building a whole new fifth generation of computers featuring AI technology; and other countries rushed to launch their own fifth generation projects. Controlling information and/or knowledge, it has been predicted, will be a key to future national survival. Information will become a commodity in and of itself. The nation that first builds and successfully applies the capabilities of 'intelligent' or super computers will have tremendous economic and political leverage.

This handbook traces the development of artificial intelligence and discusses current and future applications. The concluding chapter also contains recommended steps for organizations to follow to successfully integrate AI technology.

CHAPTER 1

BACKGROUND INFORMATION

Introduction

Artificial intelligence (AI) is considered by many to be one of the key technological innovations that will continue to shape the transformation of our society into the next century. We have moved steadily from an industrially based economy to one that is dominated by the creation and processing of information. The largest US manufacturing industry is now newspaper publishing, and the largest organization in the world in terms of employees is the Bell System.¹ Computer systems have clearly become a dominating feature of our complex society. Science fiction movies such as 2001 and War Games create the illusion that super smart computers may soon take over the world. Despite almost 30 years of research, practical intelligent systems are still in their infancy. A great deal of work must yet be accomplished if AI systems are to become commonplace. It is this ambiguity over the state of computer development that makes it difficult to separate fact from fiction.

Artificial intelligence is not clearly understood. There appears to be confusion over what AI is and what sort of functions it will be able to perform or support. The popular literature on AI is replete with promises of spectacular feats waiting in the wings. Only in the small print do we see that additional technological improvements must be accomplished first. Part of the mystique surrounding AI is that most people see it as a very complex, highly technical field of study. And AI systems are sophisticated. One need not be a computer scientist,

however, to make intelligent decisions about how to use and apply AI effectively. The first step any individual or organization should take is to become generally familiar with AI--its potential and limitations.

This handbook is offered as a starting point. It is intended as an introduction to a very complex subject. Many of the articles on AI are technical documents difficult for the layman to understand. I have attempted to provide a straightforward discussion. Artificial intelligence does appear to have tremendous possibilities for the military. This handbook is designed to help the reader determine how to begin to use this new technology.

Definitions

Is AI a field of scientific research, a technology, an advanced computer system, or what? Part of the uncertainty about AI is that there is no simple and universally accepted definition. One must be careful in discussing AI, for the frame of reference of the discussant determines what AI may mean. I have, therefore, included several working concepts of AI. You will note that some speak of AI as a field of computer science while others approach the subject as a working system. In all of these definitional discussions, it is important to realize that 'intelligence' refers to the power or act of understanding. Artificial intelligence does not refer to facts or information about a potential adversary. Avron Barr and Edward Feigenbaum's book The Handbook of Artificial Intelligence is becoming the standard reference work for AI research. Perhaps their definition is the one we should look at first.

Artificial intelligence is that part of computer science concerned with designing intelligent computer systems; that is, computer systems that exhibit the characteristics we associate with intelligence in human behavior--understanding language, learning, reasoning, solving problems, and so on.²

Patrick Winston, in his popular introductory textbook on AI, states that "AI is the study of ideas which enable computers to do the things that make people seem intelligent."³ In trying to define AI, Winston and others find themselves making frequent references to man's ability to think and act intelligently. Winston goes on to say:

AI excites people who want to uncover principles that apply to all intelligent information processors, not just those that happen to be made of wet nervous tissue instead of dry electronic gadgetry. Consequently, there is neither an obsession with mimicking human intelligence nor a prejudice against using methods that seem involved in human intelligence."⁴

Others see AI not so much a revolutionary change as an evolution in computer technology. They argue that while one goal is to make computer systems 'smarter,' they do not necessarily have to replicate the human thought process. "AI is the use of the computer to perform tasks that currently require human intelligence. AI is not a computerized re-creation of thought. This distinction between thinking like a human and performing tasks like a human is the key to understanding AI's future opportunities and limitations."⁵ Those that see AI as an improved computer tool feel that the term 'artificial intelligence' may be misleading. It conjures up images of an electronic Frankenstein. They argue, instead, that AI is merely a tool to extend the capabilities of machines. And men have been searching for ways to replicate and expand their physical and mental powers throughout history.

Origins of AI

The development of weapons and mechanical tools was simply a manifestation of man's desire to extend his own capabilities. That goal has long existed in the area of mental powers as well. Mechanical adding machines were developed as early as 1642 by Pascal. In the last century, Charles Babbage designed an 'analytical engine' that was to operate using punched cards. While his design was sound (punched cards were eventually used in calculating machines), engineering difficulties precluded successful completion of the project. Mechanical calculating machines eventually overcame the problems that stopped Babbage, however, and were in widespread use by the start of WW II. These machines were used to enable man to manipulate numbers at ever increasing speeds. In 1946, an electronic calculating machine (ENIAC) was developed to help calculate ballistics information. It was with the development of ENIAC that artificial intelligence research was born.

Artificial intelligence pioneers such as Herbert Simon and Alan Newell of Carnegie-Mellon University, Marvin Minsky of the Massachusetts Institute of Technology, and John McCarthy of Stanford realized that if computers could manipulate symbols for numbers they could also meaningfully manipulate the symbols for words, musical notes, or any other complex notion.⁶ The work of these original AI researchers did not go unnoticed.

Popular magazines such as Life, Time, and Atlantic Monthly began printing stories on the new computer age in general and "thinking machines" in particular. Predictions were made about machines that one would be able to converse with in ordinary English and not have to know

what went on inside the machine. These new marvels would be able to complete complex computations exceedingly fast and never make mistakes. A 1961 Life magazine article entitled "Machines are Taking Over" stated that computers were slowly replacing man in many endeavors but that man could always "reach down and pull the plug."⁷ Sales executives at IBM were afraid that their computers would be psychologically threatening and that customers would refuse to buy them. Ads were developed to show that computers were really pretty dumb after all.⁸

The mainstream of computer development during these early years dealt with straightforward number or data manipulation. AI pioneers, however, struggled with the concept of creating machines that could demonstrate reasoning and learning capabilities. Early projects centered on games such as checkers or chess. If a machine could be made to play chess reasonably well, it was argued, then machines could be considered intelligent. While advocates of artificial intelligence began researching 'smart' systems, others were already claiming that machines were not like humans and could never be made to 'think.'

Can machines think? No present models for understanding knowledge, how the mind works, etc., suffice to even attempt to answer this question. This introductory handbook will not attempt to dwell on a point that may be moot anyway. But it should be helpful to discuss the basic arguments surrounding this debate. Perhaps it will help remove some of the mystique that envelops AI. Hopefully it will also provide a common basis for appreciating the use of terms such as 'intelligent machine' or 'smart systems.'

Those who said machines could be made to demonstrate human-like intelligence developed the theory of the "Turing Test" to prove their point. In this test, an interrogator would be separated from the person or machine being interrogated and could communicate only by teletype. If the interrogator could not tell whether he or she was communicating with a person or machine, then a machine could be said to think. (A.M. Turing was a British AI pioneer).⁹

A number of AI researchers believe that machines do perform thinking functions. The fact that computers can't write like Shakespeare, they argue, doesn't mean they aren't intelligent. Very few humans are that creative. Does this mean the rest of humanity doesn't think either? Artificial intelligence, they stress, will "extend those human capacities we value most--properties we sum up as our intelligence or our reason; the thinking machine would amplify these qualities as other machines have amplified other capacities of our body."¹⁰ The crux of the debate over 'thinking machines' comes down to perceptions and definitions of intelligence or intelligent behavior.

Those who do not believe machines can or ever will think argue that it is impossible to replicate the human mind in a mechanical device. Computers are not innovative; they can only do what they are programmed to perform. Humans approach problem solving from a holistic sense of the world. That is, they use all their senses, knowledge, and heuristic abilities to think and respond. Computers could never achieve this level of awareness. The following is a summary of the primary arguments against the concept of thinking machines:

. . . intelligence is an exclusive human property; for reasons of divine origin or biological accident. Human beings are the only creatures on the planet who have or will ever have genuine intelligence . . . machines can't be said to think because intelligence requires creativity or originality, and no machine has been or can be creative and original. . . Given that computers might be capable of intelligent behavior ought we to pursue the possibility? Can we foresee the outcome of such an awesome step?¹¹

A leading critic of the notion that machines could think is Hubert L. Dreyfus. He attacks what he believes to be the exaggerated claims of the AI researchers. He argues that human intelligence is unique and that the general misunderstanding about AI is a result of researchers in AI stating that thinking machines are already here or just around the corner.¹² Dreyfus believes that efforts to create thinking machines are doomed to failure because information files (which replicate the human holistic view of things) would become overwhelming. Dreyfus compares AI research to a man climbing a tree to get to the moon. Progress is made in the early going, but it quickly gets harder and finally peters out in the upper branches.¹³ One of the most telling points the critics of 'thinking machines' make is the difficulty of giving computers common sense. This is a formidable problem that does not appear close to solution; but the nonbelievers have not weakened the faith of AI experts such as Patrick Winston. He stated the case in this way: "Of course to believe in human superiority is a tradition. Once our intelligence was unchallenged, yet someday computers may laugh at us and wonder if biological information processors could be really smart."¹⁴

Conventional Computers

To understand how AI systems work, let us first briefly review a few fundamentals about conventional computers. Computers, in general,

are devices that accept and manipulate data in a sequence ordered by some prearranged program. These operations result in some further action or output. Computers that perform these operations are generally divided into two basic types--analog and digital. Analog computers operate on a constant but varying input (like an automobile speedometer) while digital computers operate on inputs that are either on-off or incrementally stepped quantities represented by numerical digits.¹⁵ AI systems generally employ digital computers, so let's quickly review how they work.

Basic Computer Operations

Digital computers have three main parts or subsystems: Input/output, memory, and central processing unit (CPU). The input/output device (keyboard, monitor, printer, etc.) provides the means to enter programs and to display results. Programs and instructions are stored in the second basic component, the memory module. Interim results, computations and data are also stored in memory until they are needed for further operations. Memory modules may also use storage devices such as magnetic tape or discs. The key component of a conventional computer, the central processing unit or CPU, processes the programs or instructions in the memory module and executes the required operations. It controls the entire operation.¹⁶

All conventional computers, from the first generation machines built in the late 1940s and early 1950s through the current fourth generation systems, are essentially the same in design and operation. Generational dividing lines came about as a result of several changes in hardware technology, software developments, and operational techniques. First

generation machines, for example, used vacuum tubes, created a great deal of heat, and were very large. Second generation machines featured transistors which reduced both size and heat problems. Integrated circuit computers introduced the third generation, and very large-scale integrated (VLSI) computers initiated yet a fourth generation. Edward Feigenbaum, a leading AI expert, believes that we are currently at the end of the third generation and that fourth generation VLSI (computers) will dominate the 1980s.¹⁷

Conventional computers built during all four generations follow an operational design known as the Von Neumann process. (John Von Neumann was a computer pioneer and mathematician.) This means that computer programs are processed serially in a step-by-step operation. Each step the computer takes is spelled out in a detailed program. It can only do what it is instructed to do. A conventional computer receives data, performs simple arithmetic, and produces answers consisting of individual digits. Special programs in the computer can convert these individual digits to alphabetic characters.¹⁸ Conventional computers, then, follow rigidly formatted programs, completing one process at a time; but technological improvements have enabled conventional computers to perform these tasks at remarkable speeds. Artificial intelligence computers operate in a fundamentally different fashion.

AI Computer Systems

Artificial intelligence systems differ in both their hardware and operational programs. AI computers are built to manipulate symbols

rather than numeric values. Some of these machines are constructed to use unique AI programming languages such as LISP (List Processing language). LISP was developed by John McCarthy in 1957 for the express purpose of handling complex concepts and symbol manipulation. These special computers are made primarily by three companies: Symbolics, Inc., of Cambridge, Massachusetts; Lisp Machines, Inc., of Culver City, California; and Xerox Electro Optical Systems of Pasadena, California. A comparative look at AI systems and conventional computers should clarify how the two systems differ.

System Comparison

Conventional computers and AI systems have a number of significant differences in the way they operate. You will recall that conventional systems use primarily numeric operations, following very precise step-by-step directions. That is, to solve problems, they follow explicit algorithmic solutions. Because of the way information and instructions are structured, it can be difficult to modify or change a program. Conventional computer programs are designed to provide specific answers to a given problem. They are not designed to guess, but rather to provide solutions previously calculated and stored somewhere in the computer's memory. It was this inflexibility that led AI researchers to design machines that could simulate human performance.

Artificial intelligence systems are primarily symbolic processors. Rather than following a predefined algorithm, the AI program can more readily sort through its stored memory to determine what procedures to follow. In this approach to problem solving, solution steps are implied

but not specifically spelled out. The ability of AI systems to use 'heuristics' instead of just algorithms gives them their most unique characteristic. Heuristics have been called the art of good guessing. Heuristics enable us (or machines) to recognize promising approaches to solving problems, to break problems down into smaller problems, to overcome incomplete information, and to make educated guesses.¹⁹ It is this flexibility that enables AI systems to develop 'satisfactory' answers that may not be precisely correct but are acceptable. Another important aspect of this flexibility is the AI system's ability to 'explain' why certain decisions were made. In an AI system, the structure of the components allows the program to be more easily modified and the knowledge base updated. The table below provides a comparison of conventional and AI systems.

Table 1
AI-Conventional System Comparison
(From An Overview of Artificial Intelligence and Robotics,
NASA Technical Memorandum 85836)

<u>Artificial Intelligence</u>	<u>Conventional Computer Programming</u>
● Primarily symbolic processes	● Often primarily numeric
● Heuristic search (solution steps implicit)	● Algorithmic (solution steps explicit)
● Control structure usually separate from domain knowledge	● Information and control integrated together
● Usually easy to modify, update and enlarge	● Difficult to modify
● Some incorrect answers often tolerable	● Correct answers required
● Satisfactory answers usually acceptable	● Best possible solution usually sought

AI Research Progress

Artificial intelligence research progressed significantly in its first three decades. It has grown from a part-time pursuit of a few individuals on the fringes of computer science to a full-fledged field of study. AI researchers now hold international conferences, publish several journals, and collect a sizable share of Defense Department contract money. AI work has moved rapidly from isolated university laboratories to a number of commercial ventures. As AI research work matured, some of the early goals were modified.

An early goal of AI was to recreate human thought processes with machines. Many of the AI pioneers spent a lot of time studying the human cognitive process and then tried to build programs that could act in a similar manner. This task quickly grew too enormous for all but the most simple programs. More recent research has concentrated on solving specific problems. Leading AI experts such as Edward A. Feigenbaum now say that "it is not critical whether the methods mimic the internal structure of human behavior."²⁰ In the early years of AI research, too many promises were made for near-term major breakthroughs. When these goals were not realized, AI went through a period (the 1960s and early 1970s) of disenchantment. Recent technological developments, such as VLSI chips, have given new life to AI research. More AI tools, such as symbolic processors, are available to help design new 'intelligent' systems.

From the formative years, through the lean times and into the present period of popularity, one agency almost single-handedly ensured AI's survival. The Defense Advanced Research Projects Agency (DARPA)

supported AI research through two decades of important (and highly risky) research efforts. DARPA's steady support enabled AI to develop the fundamental knowledge and tools that are finally delivering the long-promised 'intelligent systems.' In the next chapter, we will discuss some of those systems and their potential value to the military.

NOTES

CHAPTER 1

1. John Naisbitt, "Computer Trends," Spectrum, November 1983, 110.
2. Avron Barr and Edward Feigenbaum, The Handbook of Artificial Intelligence, Vol. 1 (Stanford, Calif.: Heuris Tech Press, 1981), 3.
3. Patrick H. Winston, Artificial Intelligence (Reading, Mass.: Addison-Wesley Publishing Co., 1977), 1.
4. Ibid., 3.
5. Richard P. Ten Dyke, "Outlook on Artificial Intelligence," Auerbach Publishers Inc., undated, 1.
6. William M. Alpert, "Computers With Smarts," Barron's, 23 January 1984, 13.
7. Warren R. Young, "The Machines Are Taking Over," Life, 3 March 1961, 108-110.
8. Pamela McCorduck, Machines Who Think (San Francisco: W. H. Freeman and Co., 1971), 159.
9. Ibid., 57.
10. Ibid., 113-114.
11. Ibid., 171-172.
12. Ibid., 182.
13. John O. Green, "Making Computers Smarter," Popular Computing, January 1984, 102.
14. Winston, Artificial Intelligence, 252.
15. Roger S. Walker, Understanding Computer Science (Dallas: Texas Instruments Inc., 1981), 1-5.

16. Ibid., 1-6.
17. Edward A. Feigenbaum and Pamela McCorduck, The Fifth Generation (Reading, Mass.: Addison-Wesley Publishing Co., 1983), 17.
18. James A. Saxon and Robert E. Fritz, Beginning Programming with Ada (Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1983), 3.
19. Joel N. Shurkin, "Expert Systems: The Practical Face of Artificial Intelligence," Technology Review, Nov-Dec, 1983, 74.
20. Feigenbaum, The Fifth Generation, 77.

CHAPTER 2

ARTIFICIAL INTELLIGENCE APPLICATION

After nearly three decades of AI research, four major areas of application have emerged--Expert Systems, Natural Language, Vision Systems, and Robotics. These areas present special, and sometimes extremely difficult, problems to overcome. Progress, as a result, has not been equal across the board. Expert systems and natural language application appear to offer the most immediate possibilities for military application. They will be discussed in some detail while Vision and Robotics will be mentioned only briefly. Artificial intelligence research was initiated to extend the use of man's mental powers through the use of computer systems. It is therefore not surprising that the most mature area of AI research and application is the so-called Expert System.

Expert Systems

Expert or knowledge base systems resulted from the initial AI effort to develop computer assisted problem solving systems. AI advocates believed that computer systems could imitate intelligent human behavior. They felt confident that computer systems would be capable of solving the most complex problems. All one had to do was build a comprehensive data or knowledge base and teach the computer how to sort through the facts and come up with an acceptable answer. The first 'problems' that were tackled were games such as checkers and chess. It was assumed that if computers could be taught to solve gaming problems,

then other 'problems' could be solved. Researchers soon discovered that limiting the scope of a problem through effective search techniques was the key to building an 'intelligent' system. AI pioneers also discovered that the domain or "world" that was being investigated had to be limited for a system to provide acceptable results. A great deal of research effort in the 1960s and 1970s was invested in attempting to develop more effective search techniques. Researchers also began to focus on the data within the knowledge base of the computer. They realized that by carefully representing and limiting the domain of knowledge contained in the computer's knowledge base, they could make it easier to extract acceptable answers for a given problem. Edward Feigenbaum describes an expert system as follows:

An "expert system" is an intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution. The knowledge necessary to perform at such a level, plus the inference procedures used, can be thought of as a model of the expertise of the best practitioners of the field.¹

Expert systems are intended to provide a method of codifying certain human capabilities in an intelligent computer system. They are to not only 'magnify' human mental abilities, but perform tasks with an unerring tirelessness while serving as 'intelligent' advisors or consultants. Most expert systems are constructed similarly, and most use similar search techniques. There are, however, no specific models for designing an expert system.

System Components

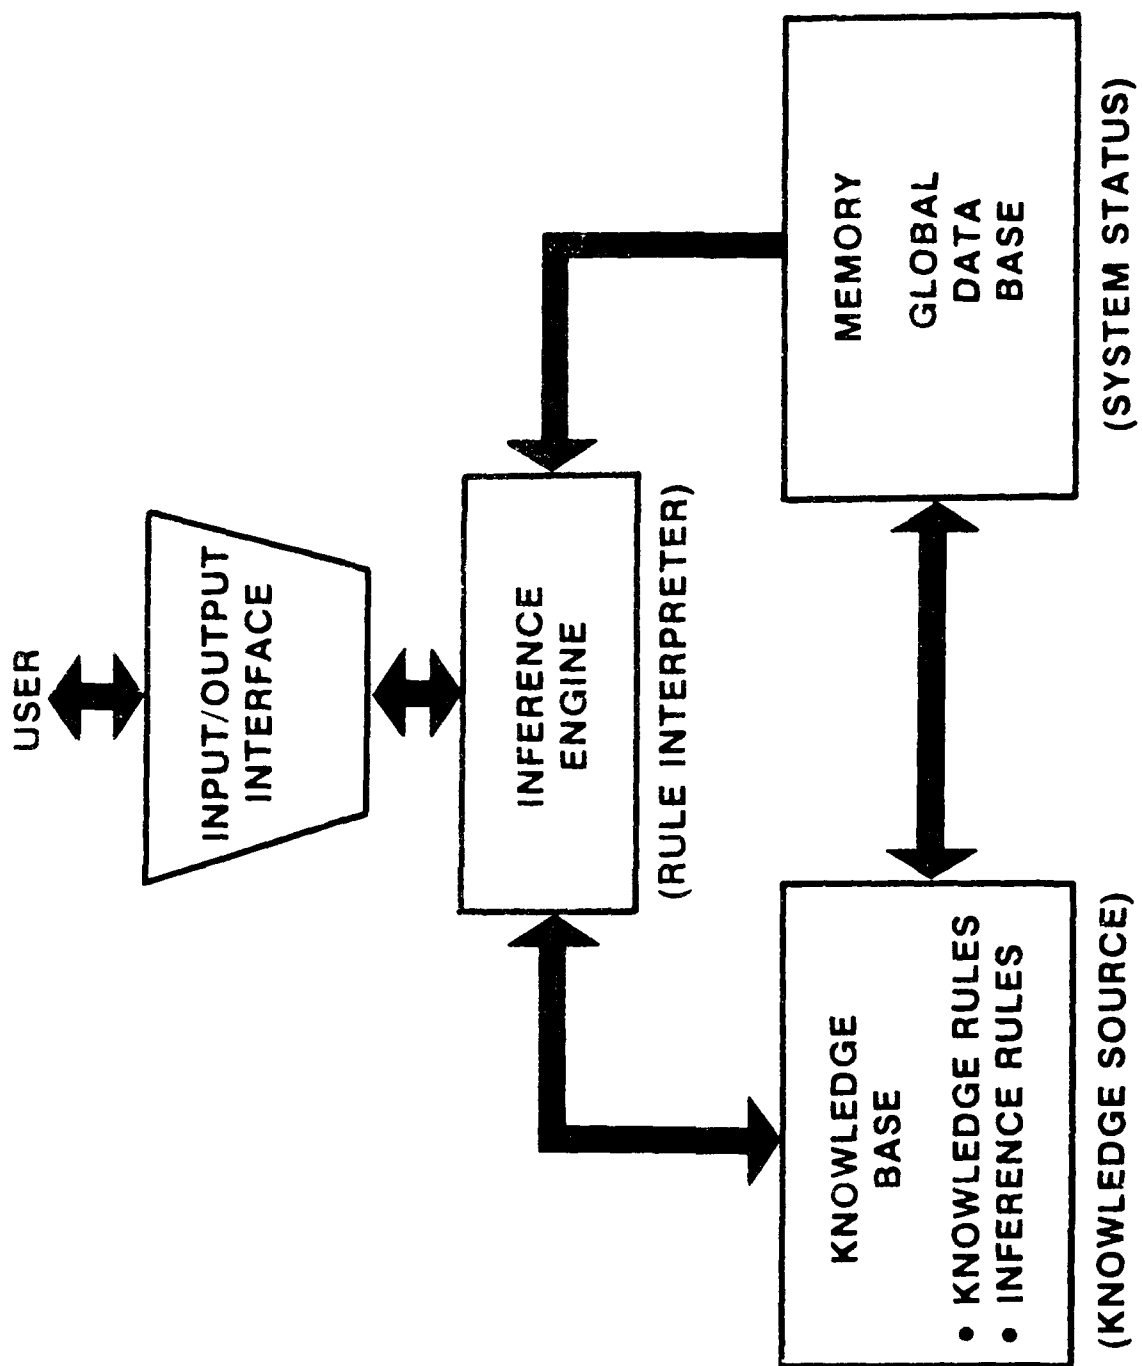
Expert system components differ from those of a conventional computer. The basic structure consists of four primary parts: an

interface device, an inference engine, a knowledge base, and a memory section (see fig. 1).

The interface device provides the user with a method of 'talking to' or instructing the system about what needs to be accomplished. It also provides the means to determine results through a video screen, printer, or other output device. Efforts to give the interface component the ability to converse in a 'natural' language are receiving increased emphasis. (Natural language usage will be discussed later in this chapter.)

The inference engine is the critical component in an expert system. It provides the control structure for sorting through the knowledge base. It determines which rules or facts should be applied to solve a given problem. The inference engine draws conclusions based on the rules in the data base. It controls program procedures and data manipulation.

The knowledge base in an expert system also has unique characteristics. Unlike conventional programs, it stores more than just facts; it also contains heuristic knowledge which replicates the expertise an expert develops in a specific domain. It is this heuristic knowledge that allows both humans and machines to make educated guesses to solve problems. Heuristics enable us to break problems down into more manageable pieces. Both the facts and the heuristic knowledge stored in the knowledge base are generally represented in the form of production rules or IF-THEN rules--"if the patient has a fever of 108 degrees, then the patient is probably very sick." An effectively constructed knowledge base is critical to the performance of an expert system. Useful knowledge an expert system might contain is discussed in Avron Barr's The Handbook of Artificial Intelligence:



Source: An Overview of Artificial Intelligence and Robotics
NASA Technical Memorandum 85836

Figure 1. Expert System

1. Facts about the domain: "The shin bone is connected to the ankle bone."
2. Hard-and-fast rules or procedures: "Always unplug the set before you stick a screwdriver into the back."
3. Problem situations and what might be good things to try to do when you are in them (heuristics): "If it won't start but you are getting a spark, check the fuel line."
4. Global strategies: differential diagnosis.
5. A "theory" of the domain: "A casual explanation of how an internal combustion engine works."

Barr goes on to say,

Much of the knowledge that characterizes human expertise is hunchlike, in the sense that it does not constitute definite consequences of actions or certainty of conclusions. . . . In particular, inexact reasoning, using hunches or heuristics to guide and focus what would otherwise be a search of an impossibly large space has resulted in systems with human-level problem solving abilities.²

Developing these special knowledge bases has been labeled "knowledge engineering." It is the job of a knowledge engineer to work with a domain expert to convert his or her knowledge into workable rules that can be manipulated by the inference engine. Knowledge engineering has become one of the critical links in developing expert systems. The procedures they follow vary and are quite time consuming. As mentioned previously, an expert system is only as good as its knowledge base. Knowledge engineers sometimes spend months developing a workable data base. It is not unusual to see over six months' research being performed before even one rule is written. Therefore, it could take years to develop a complex knowledge base! Another factor that adds to the time consuming process may be the requirement to have only one expert as a source for rules. This can preclude conflicting rules being built into a system. It is realized, however, that the sum of experience of several experts may be required for some systems.

The memory area provides a place for data to be stored while possible conclusions or solutions are being developed. This 'global data base' keeps track of the problem's status--what has already been done to solve the problem and what is needed for a particular step in the problem.

Search Procedures

The control strategy determines how the rules within the expert system are applied. There are a number of ways to control the inference engine's manipulation of the knowledge base. One of the more common procedures is known as a 'forward chaining' or data driven control strategy. In this procedure, the knowledge base is scanned to sort out those rules that are relevant to the problem or question presented. When a match is found, the rule is sent to the memory area and the search continues until no more matches can be made or until the instructions are modified to renew the search. Another search procedure is to select a desired goal and then scan the rules to find those whose consequent actions can lead to the desired state. This is known as a goal-driven or "backward-chaining control strategy."³ More complex expert systems might use a combination of search procedures. The control structure might also employ meta-rules, which would either select relevant rules from an entire set of rules or focus the search in a specific part of the data base. This helps overcome one of the key difficulties in expert systems--reducing the search space that must be considered.⁴ These search procedures can be found in virtually all expert systems now in use.

Existing Expert Systems

The first working expert systems were developed in the 1960s. Projects such as Stanford University's DENDRAL chemical analysis system and the Massachusetts Institute of Technology's MYCIN medical diagnosis consultant provided the knowledge and working tools necessary to develop other expert systems, which now number in the hundreds. Despite advances in expert system tool development, it is still a lengthy process to build an expert system. Professionals involved in the business of building these systems provide a number of cautions that prospective buyers should heed. They point out, for example, that a specific problem to be worked by an expert system must first be identified. If the problem can't be defined in detailed terms within a narrow domain, it is unlikely that a successful system can be built. Secondly, they warn that it will take more time to construct an expert system than you might think. Rule gathering and knowledge base construction can be especially time consuming. It is also advisable to build a demonstration system first (rather than initially trying to build a full scale model). Several iterations are usually required before an effective expert system is developed.⁵ The following list is a sample of some of the expert systems now in use:

SYSTEM	FUNCTION
MYCIN	Medical Diagnosis
DENDRAL	Chemical Spectral Analysis
PROSPECTOR	Geographical Analysis
R1/XCON	Computer System Configuration
AIRPLAN	Naval Aircraft Operations
HARPY	Speech Understanding

STEAMER	Steam Propulsion Plant Operation
ROSIE	Expert System Construction
BATTLE	Battlefield Weapons Assignment

This is but part of a list that is growing almost daily. New systems are being developed in almost every field where an 'intelligent' consultant would be helpful. For a more complete listing of expert systems, see William B. Gevarter's An Overview of Artificial Intelligence and Robotics, Vol 1. One of the areas of AI research and application that will help make expert systems easier to use is that of Natural Language.

Natural Language

Computers now perform amazing feats at incredible speeds, but one feature has improved very little since they were first introduced. It remains difficult to 'talk' to a computer. They do not speak the same language that you and I use, so we must learn their language. The burden of understanding between operator and computer has always been on the human. Communication with it must be very precise, and only in the programming languages it understands. If I make a mistake, even a slight one, I will not be understood. Instructions cannot be given in abstractions or analogies. Sometimes we come to believe that it is easier just to go ahead and work a problem without a computer.⁶ AI researchers are among those developing "user friendly" interfaces to change all of this.

Natural language research involves areas as diverse as translation devices, voice recognition systems, synthetic voice production devices, and interface systems that will allow operators to speak in their natural

language. Natural language programs must be flexible and have a dynamic memory in order for a system to understand within a given context.⁷ It is this latter area that is gaining the most attention in university and commercial research labs. While everyone realizes that computers are becoming an ingrained part of our social fabric, we also know there is continuing resistance or unwillingness to learn a computer language.

Natural Language Interface Systems

Researchers trying to develop computer systems that can 'understand' natural language discovered that knowledge was the key. Gevarter points out that for a computer program to interpret a relatively unrestricted natural language communication, a great deal of knowledge is required. Knowledge is needed of: (1) sentence structure; (2) word meaning; (3) word morphology; (4) sender beliefs; (5) conversation rules; and (6) an extensive shared body of general knowledge about the world.⁸

To date, machines can't interact in a very "natural" manner. To really speak in a natural language, a system should have a large vocabulary so that the operator could be understood without having to speak in a limited lexicon. An obvious reward of developing an effective natural language interface system would be to expand the powers of computing to larger portions of our society. Commercial firms, realizing the potential profit, have made this one of the most exploited AI research areas. Intellect is one such commercial program. It converts English instructions into machine language and then prints the instructions in 'natural language' on the monitor so the user knows he or she is understood. The difficulty in developing a knowledge base for

converting instructions from one language (natural) to another (artificial machine language) first became apparent during the English-Russian language translation projects of the 1950s.

Translation Projects

Translating Russian text into English was one of the first great hopes for artificial intelligence. The Defense Department's Advanced Research Projects Agency (DARPA) spent a great deal of money to make this goal a reality. The effort was such a failure, however, that it not only let back work on AI translation projects but nearly terminated all AI work. So many bold promises were made, and the results were so pathetic, that it took close to 15 years for AI researchers to come out of their laboratories again.

The first natural language processing effort . . . attempted translation by purely syntactic means: dictionary lookup of words substitution, and reordering of sentence syntax. The effort was widely viewed as a failure: "The spirit is willing but the flesh is weak" is said to have come out as "The vodka is strong but the meat is rotten" when translated into Russian and back.⁹

These early projects pointed out that mechanical translations have trouble with the fuzzy boundary between syntax and semantics. The formal structure of words in a sentence do not always explain the meaning of those words. Natural language processing (NLP) systems must use both linguistic and domain knowledge to interpret an input. These NLP systems vary in their complexity. Simple ones store facts about a limited domain and then perform pattern or key word matching. More advanced systems have enough information about the domain to understand inputs in terms of context and expectations. They might even contain information about the beliefs and intentions of communicators. Systems

that go beyond pattern matching require a more extensive language knowledge to interpret sentences. This is usually done by "parsing" a given input into component parts. AI natural language researchers have worked on the parsing problem by studying the areas of syntax (structure of phrases and sentences), semantics (meaning), and pragmatics (use of language in context). To date the only practical working natural language systems are the simple key word or pattern matching systems employing limited lexicons.¹⁰

Voice Recognition

Voice recognition systems comprise another area of natural language research that has received a lot of attention. Progress has been even slower than in the translation systems. AI researchers tend to use human capabilities as a model for their systems. Voice recognition systems follow that trend. But replicating the ability of humans to hear spoken words and translate them into meaningful messages is one of the toughest challenges AI workers will face. "Machine recognition of spoken words has been one of the problems pursued since almost the beginning of computer science, and the results have been frustratingly poor."¹¹ The Rome Air Development Center has worked for over 10 years to develop an Automatic Speech Recognition (ASR) system. The system is to be able to remove noise, identify the language being spoken, recognize important keywords in the conversation, and identify the speaker. Raj Reddy, an AI natural language researcher, explains the problem:

To understand speech, a human considers not only the specific information conveyed to the ear, but also the context in which the information is being discussed. For this reason, people can understand spoken language even when the speech signal is corrupted by noise. . . . It is difficult to develop computer programs that are sufficiently sophisticated to understand continuous speech by a random speaker. Only when programmers simplify the problem--by isolating words, limiting the vocabulary or number of speakers, or constraining the way in which sentences may be formed--is speech recognition by computer possible.¹²

Speech research is conducted in systems having various degrees of sophistication. Some systems can identify only a few isolated words spoken by a specific person. Other systems are 'speaker independent,' and may be able to recognize continuous speech. Obviously, the more complex systems present a challenge of enormous proportions. "The ability to recognize unrestricted vocabularies of continuous speech from any speaker--remains distant and certainly will not be realized until artificial intelligence progresses far beyond its current stage."¹³

A related area of research is that of synthetic voice production or speech synthesis. While it is an area of limited AI research to date, it may play an important part in 'talking' expert systems. Some applications of speech synthesis have already been introduced into automobile warning systems. Similar warning devices are being built into military systems. Computers could be given the power to provide vocal instructions, and electronic mail could be read to the receiver. Now limited, talking systems seem destined to appear throughout industry and the home.

Voice recognition applications offer one of the most rewarding natural language areas. Speech is our most natural way of communicating; but as has been pointed out, it is probably going to be the toughest challenge AI faces.

Vision Systems

Previously called 'scene analysis,' computer vision is an area of pattern recognition which is concerned with giving the computer the ability to perceive visual input data. The goal is to recognize and understand the contents of a scene through analysis of given images in much the same way that a human understands what is seen. In particular, researchers in robotics vision seek to develop a visual sensor technology for "industrial robots and human-like manipulators for operation in an unpredictable environment."¹⁴

Computer vision is still in the early stages of development. Human sensory capabilities provide the model to be emulated as they do in other AI projects. Computer vision involves many different processes which integrate different kinds of information. The Handbook of Artificial Intelligence divides computer vision into signal processing, classification, and image understanding. Before analysis can be made, an image must be described for a sensory input. The first stage of the general analysis process involves feature extraction. Based on visual data, the main features of an image are marked or 'extracted.' The second stage involves segmentation. The third stage is property measurement. Measurable properties (e.g., range, area, texture) of these image segments are used to obtain a better description. Finally, stored models (general descriptions of predefined objects) are compared with descriptions inferred from current processing. Matching results in an identification with the appropriate properties being associated with the original image.¹⁵

Computer vision applications are found in medicine, remote sensing, and industry. Reading textual data is an old application of vision which involves character recognition. Handwritten text is obviously more difficult to 'read' than machine generated text. Bar codes, printed on almost all commercial packaging, are an example of a character recognition system. The cash register 'reads' the code to determine a price. Computer vision is applied in tissue analysis to perform cancer or blood cell counting. The earth's natural resources can be monitored using images provided from remote airborne sensors. Applications in industry include robotic vision for locating and removing defective products. Considerations in designing a robotic vision system include cost, real-time operation, reliability, and flexibility.¹⁶

While computer vision research has progressed considerably in 25 years, several problems persist. Successful applications to date involve images with simple domains that are primarily two-dimensional. Three-dimensional images are more complex to analyze and process. Recognizing non-overlapping parts on a conveyor belt under good lighting, for example, is already being accomplished. Recognizing overlapping parts in a shadowy bin poses very difficult problems. A two-dimensional image of a three-dimensional object does not provide enough information. Secondly, the influence of ambient factors on sensor perception is unknown. Another problem is that the vision process is not well understood. How human vision is actually accomplished is unknown. The problem is that of trying to give the computer the same holistic sense of its surroundings that humans with

normal vision perceive. An AI vision researcher stated: "The development of a general-purpose computer vision system that can approach the abilities of the human eye and brain is remote at present, despite recent progress in understanding the nature of vision. . . . Clearly the general problem of recognizing objects in a scene and describing their relations in three dimensions is far from solved. Existing systems can deal only with restricted types of scenes and they operate slowly."¹⁷ Researchers at the Army's Topographic and Engineering Laboratory have been working since the 1960s on a system that could extract information from aerial photography using pattern recognition. Only limited progress has been made. DARPA has also contracted with various universities to investigate methods and techniques to develop automated machine capabilities. Gevarter notes that "the amount of activity and the many researchers in the computer vision field suggest that within the next 5 to 10 years, we should see some startling advances in practical computer vision, though the availability of practical general vision systems still remains a long way off. . . ."¹⁸ Computer vision developments are especially important when combined with advances in robotics.

Robotics

Advanced automated military systems will depend on progress in the field of Robotics. Systems that can operate virtually independently have tremendous potential in the military, as well as in the industrial sector.

The coupling of automatic production machinery . . . is a primary motivation for increased use of sophisticated, programmable robots. A robot that can sense its surroundings,

modify its actions, carry out specified tasks and be given tasks on short notice in languages natural for the task is exhibiting rudimentary intelligent behavior. Ultimately, it is this exhibition of intelligent behavior that will distinguish robots from increasingly sophisticated conventional machinery.¹⁹

Robotics have already been used to explore the surface of other planets, and are presently used in a number of manufacturing areas. Mobile Robots of Woburn, Massachusetts, for example, has developed robots for prison application. Using ultrasonic and infrared sensors, robots will be used as guards to detect human odors and participate in riot control.²⁰ To improve industrial robots, researchers are working to combine improved robotics with improved computer vision systems. The result should provide systems that can 'see' what they are doing and respond more 'intelligently.' Military planners envision using such systems as automated reconnaissance vehicles or automated aircraft. Before such systems can be successfully deployed, a great deal of additional research in both robotics and computer vision will be necessary.

Development Problems

Artificial intelligence research promises great things for the future, but it still faces a number of significant problems. Some of the technological difficulties, such as in speech recognition and effective visual systems, have already been noted. These technical problems are discussed in greater detail in the many AI publications now available. Other difficulties range from misinformation to overzealous claims.

Although AI has been researched for almost three decades, the number of AI experts remains relatively small. There are, for example,

few knowledge engineers available to convert 'expert knowledge' into quality data bases for expert systems. While new AI development tools are becoming available, a lot of knowledge engineering is still done by hand. This is likely to continue to be so into the foreseeable future.²¹ There is a shortage of AI researchers in general, which is being exacerbated by the rush of qualified personnel into commercial ventures. Difficulties in transporting software programs from facility to facility is another major problem AI researchers face because of hardware differences.

Artificial intelligence programs are written in a special language developed specifically for manipulating symbols rather than numeric values. This language, called LISP for list processing, was developed by John McCarthy in 1957. It has become the standard programming language for most AI researchers, though the Japanese and some Europeans use another AI language called PROLOG. LISP allows AI systems to be developed in a more flexible format than conventional programming languages. It is this programming language difference that could pose problems for AI application. First of all, programs written in LISP may not be portable from one computer system to another. Secondly, military computer programs are written in a rigidly formatted language (such as COBOL). Complicating the picture even more may be the fact that the Department of Defense (DOD) has adopted Ada (another structured language) as the official program language for embedded computer systems (missile guidance, for example). If AI systems are to be generally accepted for military application, computer hardware and software compatibility problems will have to be resolved.

Following the translation failures of the 1950s, AI advocates became more cautious in their claims for the future. For a period of almost 15 years, AI endeavors went virtually unnoticed. In the late 1970s, this all began to change. Technological improvements, such as VLSI chips, had given AI researchers some of the tools they needed to begin making progress once more, and society in general was experiencing a computer revolution. By 1983, AI had become a hot area for venture capital once more. This turn-around in AI's fortunes had mixed results. Many AI researchers feel that "public expectations for AI have gotten dangerously overheated, with the concomitant risk of disappointment and backlash . . . the hoopla has left researchers bemused and concerned . . . the promises seem to be outracing the reality."²² Companies are rushing to cash in on AI's new-found popularity. Business Week reported that

with nearly 40 small companies vying for a place in the AI market, competition is intense. And some companies have already gotten into trouble in their rush to bring projects to the market . . . experts fear an "overselling" of technology . . . Without question, some of the AI products now entering the market are not derived from AI technology at all. Some companies openly admit that they have simply relabeled existing software to cash in on the AI boom.²³

The point of all this is that we must be aware not only of the major technical problems that must be solved before effective AI systems can be deployed, but the misinformation we may hear as well. Despite these various difficulties, AI does appear to offer tremendous potential for military applications in the future. In the next chapter, we will look at where AI research is headed and how this technology might be used successfully.

NOTES

CHAPTER 2

1. Edward Feigenbaum, "Knowledge Engineering for the 1980's," Computer Science Department, Stanford University, 1982, 1.
2. Avron Barr and Edward Feigenbaum, The Handbook of Artificial Intelligence, Vol. 1 (Stanford, Calif.: Heuris Tech Press, 1981), 81.
3. Richard O. Duda and John G. Gaschnig, "Knowledge Based Expert System Come of Age," BYTE Publications Inc., September 1981, 246.
4. William B. Gevarter, An Overview of Artificial Intelligence and Robotics, Vol. 1-Artificial Intelligence, Part A-The Core Ingredients (Washington, D.C.: National Aeronautics and Space Administration, 1983), 28.
5. Harvey S. Goodman, Head, Software Applications Section, TRW Systems Group, Seminar on Expert Systems, Orlando, Florida, 15-16 October 1984.
6. Steven K. Roberts, "Artificial Intelligence," BYTE Publications Inc., September, 1981, 178.
7. Robert C. Schank and Peter G. Childers "Experiments in Artificial Intelligence," Computerworld, 12 November 1984, ID/1-ID/28.
8. William B. Gevarter, An Overview of Artificial Intelligence and Robotics, Vol 1-Artificial Intelligence, Part B-Applications (Washington, D.C.: National Aeronautics and Space Administration, 1983), 41.
9. David L. Waltz "Helping Computers Understand Natural Languages," Spectrum, November 1983, 82.

10. Gevarter, An Overview of Artificial Intelligence and Robotics, Part B, 42-52, 10.
11. Bertram Raphael, The Thinking Computer (San Francisco: W. H. Freeman Co., 1976), 206.
12. Raj Reddy and Victor Zu, "Recognizing Continuous Speech Remains an Elusive Goal," Spectrum, November 1983, 84.
13. Harold L. Andrews, "Speech Processing," Computer, October 1984, 318.
14. K. S. Fu and Azriel Rosenfield, "Pattern Recognition and Computer Vision," Computer, October 1984, 274-282.
15. Azriel Rosenfield, "Why Computers Can't See (Yet)," Abacus, Vol. 1, 1983, 17-26.
16. Kevin McKean, "Computers That See," Discover, September 1984, 72-74.
17. Takeo Kanade and Raj Reddy, "Computer Vision: The Challenge of Imperfect Inputs," Spectrum, November 1983, 88.
18. Gevarter, An Overview of Artificial Intelligence and Robotics, Part B, 39.
19. John S. Jarvis, "Robotics," Computer, October 1984, 283.
20. Barbara Wierzbicki, "Mobile Robots Scheduled for Testing as Prison Guards," Infoworld, 20 February 1984, 17.
21. Frederick Hayes-Roth, "Codifying Human Knowledge for Machine Reading," Spectrum, November 1983, 81.
22. M. Mitchell Waldrop, "Artificial Intelligence (I): Into the World," Science, Vol 223, 24 February 1984, 802, 804.
23. Business Week, "Artificial Intelligence is Here," 9 July 1984, 61.

CHAPTER 3

FUTURE PROSPECTS

Overview

Computer technology has developed at an incredible pace, and the world is transitioning into a society that literally lives on information. In the past, power has been measured by such elements as territory controlled, annual production output, and so forth. A new basic element of power may have to be added to the list. In the future, nations that control information or knowledge may possess a major source of influence in international affairs. The systems which may make such control possible are to be the products of the so-called fifth generation of computer technology. These new systems will represent a distinct break with conventional von Neumann type computers. Parallel or concurrent architecture will allow machines to conduct a multitude of operations at the same time. Advanced software designs, VLSI, and artificial intelligence technologies will give fifth generation computers expansive capabilities.

One of the primary requirements of an effective AI system is rapid access to a vast memory capability. New microelectronic developments in areas such as VLSI chips are providing that key ingredient. Another factor that will contribute to AI's probable expanding role in future technologies is its improving record of performance. Hundreds of expert or knowledge base systems are being developed. Artificial intelligence technology is now commercially profitable. A compelling factor (and perhaps the most vital) that dictates investigation and support of

future AI research and application is the edge it may provide in coping with the future. Robert E. Kahn of the Defense Advanced Research Project Agency (DARPA) summarized the importance of the next generation of computers.

We are now at a stage where the confluence of these two disciplines--microelectronics and artificial intelligence--may indeed produce new generations of computers that are both fast and smart. . . . The nation that dominates this information-processing field will possess the keys to world leadership in the twenty-first century. . . . Better planning coupled with effective execution of plans can be decisive in maintaining national security.¹

Leading computer scientist and AI pioneer Edward Fiegenbaum supports this position: "The world is entering a new period. The wealth of nations, which depend upon land, labor, and capital during agricultural, and industrial phases--depend upon natural resources, the accumulation of money, and even upon weaponry--will come in the future to depend upon information, knowledge, and intelligence."² Information in the future will become a commodity in and of itself, and the nation that fails to prepare for this transformation may be doomed to a second class status. New fifth generation computing machines, featuring artificial intelligence will give man the ability to amplify his 'knowledge' power.

Conventional computers featured von Neumann's serial or step by step processing of programs, but fifth generation computers will be built differently. New programming languages and the capability to handle symbols as well as numbers will be developed. Machines will be designed to manipulate data 'intelligently' rather than performing mere data processing.³ Artificial intelligence will provide the cornerstone for building these new systems. Robert Kahn described how these systems

might look: "The new generation of machines will consist of collections of these (AI) systems integrated into a whole. If a system needs vision, it can be added. Speech input can also be incorporated. . . ."4 Fifth generation computers are expected to play an increasingly important role in the management of industrial and military resources. The obvious advantages and benefits to be gained from AI technologies caused a number of countries to make AI the centerpiece of new computer development programs.

Fifth Generation Projects

During the early 1980s, a number of countries came to realize the potential of evolving computer technology. Japan, the United States, Great Britain, the European Common Market, and the Soviet Union all embarked on massive new computer projects to ensure that they weren't left behind in a computer revolution. Goals of the individual nations, as we shall see, differ slightly, but their overall objectives and approaches are similar. Japan, many observers believe, was responsible for this new emphasis in computer research. They shocked the world when they announced their plans for the future at the International Conference of Fifth-Generation Computer Systems held in Tokyo in October of 1981.

Japan's Fifth Generation Computer Project

Japan's industrial power was built on its ability to successfully apply western technology to Japanese industrial programs. Video recording technology, for example, was first developed in the United States; but virtually all video recorders now sold are produced in

Japan. Their success in surpassing America in auto production is also well documented. The Japanese believe that the computer industry holds the key to continued economic success. Japan made it clear at the 1981 International Conference that it would no longer be satisfied with just developing other nations' technologies. They set forth a national goal of becoming number one in the computer industry.⁵ Fiengenbaum attended the conference and was immediately struck by the significance of their plan.

They aim not only to dominate traditional forms of the computer industry but to establish a "knowledge industry" in which knowledge itself will be a salable commodity like food and oil. The (Japanese) plan documents a carefully staged ten-year research and development program on Knowledge Information and Processing Systems . . . the Japanese understand that if they succeed in this visionary project, they will acquire leverage over all kinds of industries, at home and abroad. The Fifth Generation (plan) is an exquisite piece of economic strategy.⁶

The Japanese Fifth Generation Plan is designed not only to give them a worldwide edge in computer technology, but to help with a number of domestic goals as well. They expect to provide better social services such as improved (and more personal) health care by using expert systems as consultants. Manufacturing industries will gain through the more 'intelligent' use of scarce energy resources. Fifth generation technology is also expected to improve agriculture output by providing more accurate weather forecasts and an improved distribution system. To achieve these goals, Japan broke with traditional economic patterns--it formed a consortium of eight manufacturers.

The Institute for New Generation Computer Technology (ICOT) was launched in 1982. The eight companies participating in the project provided both the capital and the personnel necessary for the venture. It is interesting to note that, except for a few research leaders, the men chosen for the project are all under age 35. ICOT's director, Kazuhiro Fuchi believes that only young men demonstrate the required zeal and innovative thinking required of this unique project.⁷ The center's ten-year plan involves three stages: (1) reviewing and evaluating current research on knowledge processing; (2) establishing subsystems for hardware and software; and (3) integrating software and hardware subsystems to create the first fifth generation computers. Japanese researchers, unlike their US counterparts, chose PROLOG (Programming in Logic) as their language for AI research (rather than LISP). PROLOG facilitates processing knowledge in parallel as well as in sequence, and it computes very rapidly.⁸ The Japanese are aware that US LISP systems will not be compatible with their more advanced machines. The significance of Japan's Fifth Generation Plan sent reverberations around the world. After a few false starts, other nations began responding to the challenge.

US Fifth Generation Projects

The United States enjoys a lead in computer technology as a result of research conducted over the years by universities, corporations, and the Department of Defense. In his book The Fifth Generation, Edward Feigenbaum warned that if the United States did not respond to the Japanese challenge its precious technological lead would rapidly melt away. "Our business as usual attitude, our near-term R&D planning

horizons, our fratricidal competitive zeal and paranoia over proprietary rights, and our planning vacuum at the national level are causing us to squander our valuable lead (in computer technology) at the rate of one day per day."⁹

It was not until the significance of the Japanese Fifth Generation Plan was finally grasped that coordinated US research efforts began to emerge. Another factor contributing to this new willingness to cooperate in joint ventures was the rising cost of conducting research. United States efforts can be generally divided into governmental programs and commercial enterprises.

Government Projects

DARPA's Strategic Computing program is the most significant US government response to the Japanese initiative. It would be incorrect to imply that DARPA was not heavily involved in computer research and developments long before the Japanese announced their plans; but partially because of Japan's initiative, DARPA has scheduled \$1 billion in new studies over the next decade. Significant sums will be invested in research efforts involving artificial intelligence, software improvements, and computer architecture. It is hoped that a new generation of 'intelligent' computers will be available for various military applications. DARPA believes that these projects will not only improve military systems but will help strengthen the US industrial base as well.¹⁰ DARPA outlined its plans for artificial intelligence as follows:

In artificial intelligence applications, software will be developed in six generic areas that hold the key to military applications: speech, vision, natural language, very large

knowledge bases, graphics, and navigation. . . . The goal is to develop real-time processing systems that can cope with the enormous computational requirements needed for exploiting these technology areas.¹¹

Universities remain a primary source for basic research work on artificial intelligence. Their research centers are receiving aid from a number of 'cooperatives.' The Semiconductor Research Cooperative, for example, is helping to solicit funds to offset the tremendous expense of performing basic semiconductor research. The Microelectronics Center of North Carolina directs funds into five different universities. Almost every major electronic firm engaged in artificial intelligence research has a working arrangement with one or more universities. This enables both organizations to reduce personnel costs and to provide for the joint use of expensive equipment. This industry-university cooperative effort is not just occurring at major universities. "Today more and more schools are integrating industry research with teaching to give students a more practical learning experience . . . universities are taking the lead in establishing industrial parks to foster joint university-corporate research."¹²

Artificial intelligence research is also being accelerated by a number of other government agencies. The Army is conducting AI research at both its Signals Warfare Laboratory at Vint Hill Farms Station, Virginia, as well as at its Engineering Topographic Laboratories at Fort Belvoir, Virginia. The Navy is working on AI decision aids and robotics at its Office of Naval Research Laboratories in Arlington, Virginia. The Rome Air Development Center has been performing AI research for the Air Force for over two decades. A new center for AI

studies is being established at the Air Force Institute for Technology (AFIT). In addition to these and other government centers for AI research, several commercial firms are entering the AI development arena.

Commercial AI Projects

Austin, Texas, edged out more than 50 other cities to become the new home of the Microelectronics and Computer Technology Corporation (MCC). The enthusiastic campaign that enabled Austin to bring this high-technology future to Texas included pledges of monetary support to build facilities and improve local university programs in computer science. Retired Admiral Bob R. Inman, former deputy director of the Central Intelligence Agency, was selected to become the new firm's chairman. He unequivocally states that MCC is a direct response to the Japanese Fifth Generation program. MCC is unique in the American business world for it brings together 19 different companies whose normal mode of operation is good, old-fashioned competition. In this venture, each company will be a cooperative shareholder in one or more of MCC's four research and development programs. An MCC promotional brochure explains the operational concept:

The concept was to join together talent and dollars to conduct pro-competitive, long-range research aimed at achieving significant advances in microelectronics and computer technology. Each company would then use the technology to develop its own products, and compete with those products in the marketplace. This concept of sharing advanced research has long been alien to American industry. As the technology dependent industry matures, however, the high costs and risks associated with basic R&D on advanced technologies have made a comprehensive research program by a single organization virtually impossible.¹³

MCC divided its ambitious project into four major programs designed to result in 'significant progress' in 5 to 10 years. Programs include packaging (improving semiconductor packaging and interconnect technology), software technology (improving the software development process), VLSI/computer-aided design (CAD) (improving technology), and Advanced Computer Architecture (the most ambitious of the proposed programs). Artificial intelligence relates to research in all the areas, but it is a specific project of the advanced architecture program. MCC hopes to improve the computer's problem-solving capabilities by developing better ways to represent human knowledge and apply human expertise to various problems. The directors of MCC realize that they must overcome a number of challenges, and not all of them are technical, before their programs become a reality. There are, for example, proprietary difficulties. Since not all of the participating companies are shareholders in all of the four major programs, transfer of technological advancements from one area to another may be difficult. Despite such difficulties, MCC represents corporate America's primary effort to preserve its lead in computer technology. And it is not just the Japanese that concern MCC now; other countries are also developing high technology computer research programs.

Great Britain

The British team of government officials and academicians that attended the 1981 International Conference in Tokyo came home convinced that the Japanese Fifth Generation Plan could not go unchallenged. Britain responded with its own five-year plan designed to improve research in four areas: very large scale integration (VLSI), software

engineering, intelligent knowledge-based systems, and man-machine interfaces. The program's expected cost of over \$500 million will be raised from both public and private sources. Members of the British computer industry are convinced that this cooperative program is necessary to improve their position in the world information-technology market.¹⁴ While developing their own national program, the British are also supporting Common Market attempts to improve Europe's competitiveness with the Japanese and the Americans.

Common Market Project

European companies hoping to improve their high tech profit share launched the European Strategic Programme on Research in Information Technology (Esprit). This joint plan calls for over 2,000 researchers to work in five primary areas over the next five years: advanced microelectronics, software technology, advanced information processing, office automation, and computer-integrated flexible manufacturing. The directors of this project realize that if they are not successful, Europeans may have to abandon the lucrative high technology market. Esprit's members believe that the information processing industry will be an especially strong growth area. They will therefore be concentrating most of their funds and efforts to make more user-friendly knowledge-processing systems.¹⁵

Soviet Fifth Generation Project

Common Market countries will find more than Western competition facing them in the future. The Soviet Union and its allies have also entered the fifth generation race. The Soviets, in cooperation with

six Eastern European partners, have launched a five-year computing plan. The Moscow Academy of Sciences will coordinate the project, supported by over \$100 million in state funds. The Soviets claim that the program is to be a civilian initiative and that they seek cooperation and help from other nations. One Soviet spokesman stated in an Izvestiya article that "developing the next generation of computers was of paramount international importance--close to another space or missile race." The Soviet program is to cover five strategic areas: design and manufacture of VLSI microprocessors, development of parallel and multiprocessor architectures, design of operating systems to better support logic programming, creation of problem-solving software, and development of expert systems and user-responsive applications.¹⁶ Kremlinologist opinions vary on how successful the Soviets might be in this project. Some argue that their innovative ability in the high tech area is too limited and that the \$100 million effort is too modest. Others argue that the Soviets are already producing high quality random access memory (RAM) chips. The US embargo on electronic technology, the Soviets claim, will not slow their progress. Soviet leaders state that just as they overcame the more complex problems of building long-range missiles so shall they overcome this latest technological challenge.

Potential Applications for AI Technology

For some time, hoped-for AI miracles have always been just over the horizon. No great technological breakthrough has occurred yet, but substantive progress is moving practical AI systems onto this side of the hill. Expectations for fifth generation AI systems abound in both

civilian and military sectors. Feigenbaum describes the potential impact of AI systems:

. . . the fifth generation (of computers) can improve and streamline medical and related information processing systems for health management, help develop systems for enabling the physically handicapped to become active, contribute computer-aided instruction systems for the lifetime education of the aged, and develop distributed processing systems for enabling people to work at home . . . decision support systems will provide high-level information for increasing the effectiveness and reducing the time and costs required for making decisions. . . . With these achievements, activities in all facets of society will be affected and within a margin of safety, more advanced, humane behavior will be possible.¹⁷

The home personal computer boom is expected to continue. Many experts believe that this trend, in conjunction with AI developments, will put 'expert consultants,' such as legal and financial advisors, in the home. AI natural language systems are expected to make computers easy to use by the masses. People will be able to give simple verbal instructions to their computer. AI machines will become even 'smarter' as they attain the ability to 'learn' on their own from various sources.

Military planners hope that AI machines and technology will help them cope with the informational deluge that is swamping decision makers. AI systems are needed that can support command and control functions, sort through and evaluate reams of data, and suggest possible actions. AI technology will be put to work on such difficult problems as command and control, data fusion, intelligence analysis, and various robotics applications. A Business Week article stated: "The military, too, is hoping that AI will eventually play an important role in the future. A 1981 study by the Defense Science Board ranked AI in the top 10 technologies for the 1980s. . . . Eventually military planners

expect machines to serve as equal partners with humans in decision-making."¹⁸

Projects such as these may sound farfetched to pragmatists. They want to see hard evidence that AI technology isn't just a pipe dream still years away from application. As previously mentioned, DARPA and a number of commercial firms are investing sizeable chunks of money, time, and personnel to prove that AI technologies can be put to use in the near term. The following sections provide examples of AI systems, both military and civilian, that are under development.

Military Applications

Expert systems represent the most mature area of AI research. A number of institutions are now working to apply that experience to military support systems.

TRW, as part of its AirLand Battle 2000 study, is reviewing the possible uses of AI technology. Intelligent machines would collect data from a number of sensors and then help formulate military responses. "Such a system would operate on an AI data base that would rapidly draw out and prioritize relevant intelligent data collected from several sources."¹⁹

Stanford Research Institute's (SRI) Artificial Intelligence Center is applying AI techniques to the problem of integrating information from various sources to help evaluate and understand developing threat situations. The goal of their work is to compile air defense order of battle, determine the identity of threats, locate the threats, and determine the status of the threat.²⁰

The US Army Signals Warfare Laboratory is conducting research to develop an automated, artificial intelligence-based tactical situation/threat assessment test model. The Army's Engineering Topographic Laboratory is also researching the applicability of AI technology to improving the radio analysis process.

The AIRFIRE Corporation, under the sponsorship of the Rome Air Development Center (RADC), is developing an advanced decision support system (KNOBS). KNOBS is a demonstration system to support the planning of counterair missions.

BRCSOFT, Inc., is developing a prototype expert system to aid the decision-making associated with tactical command, control, and communications countermeasures (C³CM). The system is to help analysts determine the best countermeasures to use against critical targets; that is, to destroy, disrupt (jam), and so forth.

Other companies and agencies are investigating the use of AI technologies to help fuse intelligence data. We have become very good at collecting data. What is needed now are systems that can provide assistance in making effective use of that information. Intelligence fusion is expected to be one of the primary AI application areas. Expert systems will help sort through information, evaluate data, and suggest possible responses.

Natural Language

Natural language systems, designed to make computers easier to use, will also be major contributors to AI technologies in the future. Conventional computers require that operators learn a special language.

The goal of researchers working on natural language systems is to put the burden of learning on the computer. Natural language input/output systems can improve the man-machine interface. Rome Air Development Center (RADC) is probably the leading research center for automatic speech recognition (ASR). RADC is working on systems that can understand and respond to verbal communications. The advantages of being able to 'talk' to a support system rather than punching buttons is obvious. It could free a busy pilot, for example, to perform other cockpit functions. ASR project directors feel confident that this technology is a near term reality. RADC developed an ASR system for cartographers at the Defense Mapping Agency. Previously, data about geographic features had to be entered in a slow manual process that diverted attention from the maps being worked on. Now, information is given verbally to a computer that 'understands' a limited vocabulary. And RADC has been working for over a decade to develop an ASR system that can listen to a radio broadcast, remove the noise, identify the language spoken, understand keywords, and identify the speaker.²¹

The Navy is developing a system which will aid maintenance personnel in diagnosing avionic equipment. Operators will be able to make verbal queries about maintenance procedures and receive spoken responses. This Aircraft Speech Interviewer will allow operators to focus their attention on repairing equipment rather than flipping through manuals. The Naval Postgraduate School is attempting to use ASR to interact with the Advanced Research Project Agency Network by voice command rather than with traditional 'keyboard' languages.

The Office of Naval Research is doing research work on applying robotics to various Naval requirements. Research activity in the future will center on new robotics software featuring AI technology. Applications will include robotic systems that can perform underwater assembly and maintenance mine sweeping, search and recovery, and scientific missions.²²

These are but a few of the AI projects presently envisioned by military planners. If the DARPA strategic computing program is successful, a "whole new generation of computers with capabilities including vision, comprehension of speech, and reasoning" will be created.²³

Civilian Applications

Commercial application of AI technologies is already well underway. Hundreds of expert systems are now on the market. Business Week stated that "optimistic analysts are predicting that AI will become a multibillion dollar annual business well within a decade."²⁴ AI is being used to automate industry and offices. It is providing advice on health care and where to drill the next oil well. Civilian application of AI in the near term will be primarily in the areas of expert systems and natural language. The following examples should underscore the fact that AI systems are a proven technology.

OPGEN is a knowledge-based system used by Hazeltine Corporation to plan the sequence of procedures to follow in the assembly of printed circuit boards. It produces, in 90 seconds, operational instructions that it takes industrial engineers 20 hours to produce.

The French oil company Elf Aquitaine will soon begin using an expert system, DRILLING ADVISOR, to help them avoid accidents that might occur during drilling operations.

General Electric will soon be using expert systems to help maintain the locomotives they build. This troubleshooting system incorporates the 'expert knowledge' of GE's top locomotive field service engineer. It will walk an engineer, even a novice one, through required repair procedures using drawings and 'how to' instructional demonstrations.²⁵

Intelligenetics' KEE and Xerox's LOOPS are new LISP software 'tools' designed to help others build expert systems. And Teknowledge has introduced M.I as an AI development tool to be operated on an IBM PC. It is designed to allow novices to create knowledge-based systems.

CLOUT is a natural language interface that allows users to converse in their native language. The system combines the capabilities of natural language technologies with that of an expert system's heuristics. Users will be able to generate complex reports with just a few questions posed in plain English.²⁶

The University of Texas is experimenting with ISAAC, a natural language-understanding computer program capable of comprehending connected discourse at 5,000 words per minute. The program is designed to help codify the rules that humans use to solve problems. "Isaac provides both a model of the processes human subjects use to solve the (physics) problems and a theory of how physical representations enter into these processes."²⁷

Infotym is introducing a new decision support system, REVEAL, which will allow users to build models with English-like statements. The

system will incorporate expert system technology to improve and simplify the use of decision support systems.

Transformational Question-Answering (TQA) System is an experimental system which is an outgrowth of early work in machine language translation projects at IBM. "TQA is focused on techniques for querying relational data bases in English, in particular, IBM's SQL (Structured Query Language) relational data base system. TQA translates English-language questions into a formal data base language, allowing users to interact with computers without being required to use a programming language."²⁸

Artificial intelligence technology is here, and it is going to play an important part in reshaping our technological society. Super computers of the fifth generation will be used for forecasting the weather, designing aircraft and automobiles, exploring for minerals, and improving the health and welfare of individuals. A Newsweek cover story stated:

. . . current supercomputers are only at the threshold of what computer designers think can be achieved; the next generation of advanced supercomputers will make today's machines look like handheld calculators. . . . The great danger for the losers in the race (to build fifth generation supercomputers)--and the opportunity for the winner--is that whoever builds the next generation of computers will have a huge technological and commercial advantage: these computers will be used for computer and microelectronics design--to build even smarter and more powerful machines.²⁹

Summary and Conclusions

When new technologies come along, there is always a tendency to either overrate or underrate their potential. Long-range bombers, for example, were going to make it possible to win wars without fighting major land battles. Computers were going to free mankind from the drudgery of processing information. Fifth generation computers, using

AI techniques, are being promoted as the new solution to many of society's problems. AI systems should, if developed properly, prove to be valuable technological tools. It must be realized, however, that AI is merely another problem-solving technology. Problems that cannot be clearly defined and/or identified using conventional methods are not going to be magically solved using AI. Artificial intelligence provides an alternate or supplementary means to work complicated tasks. It is not going to replace or do away with existent computer support systems. That is why some computer scientists view AI as an evolutionary extension of present computer capabilities rather than as a revolutionary technology.

What does all of this portend? A number of points should be considered. First, AI is not a chimerical dream to be explained by using smoke and mirrors. It is in fact a sound technology built on nearly 30 years of intensive research. But because AI involves building 'thinking' machines, it is both misunderstood and misrepresented. AI systems capable of performing the science fiction feats of a HAL (from 2001) or a WOPR (War Games) may never be constructed. Systems that can function 'intelligently' as consultants or advisors will, however, become commonplace in the next decade. That brings us to the next point.

The military must not ignore the tremendous potential offered by AI supported systems. Government agencies such as DARPA and RADC fully grasp AI's significance and are pursuing its development. For AI to be widely accepted (and implemented), however, senior military leaders will need to learn more about it. They need to know enough about AI to separate fact from fiction. Only by becoming more familiar with this

technology can commanders effectively apply AI to their own organizational requirements.

Those presently working in AI research and development will tell you that AI is not something you can buy off the shelf like paperclips. Before any organization invests time and money in AI technology, the following should be considered.

- Learn about AI through a literature review. A list of recommended documents is at attachment 1.
- Establish an AI steering group or focal point that can work with other organizations.
- Visit governmental and private AI research centers.
- Review organizational needs to select candidate areas for AI application. The organization's AI steering group should oversee this operation.
- Select a likely demonstration project to work on a real, CLEARLY DEFINED problem. Remember that it may take months or years to develop a rule-based expert system.

Throughout this process, organizations should closely monitor the progress of other AI research endeavors. Computer scientists frequently remind us that real progress in AI is going to depend upon a cooperative effort by all of those working in the field.

When this project began, I was admittedly skeptical about the potential of artificial intelligence. Too many times in the past, computer systems have failed to perform as advertised. Even the Japanese are beginning to hedge their bets on their Fifth Generation

project. At the Second International Conference on Fifth Generation Computers, the Japanese revealed that they were modifying goals and pushing target dates into the future. Some Japanese ICOT firms, believing the Fifth Generation Plan unrealistic, are quietly pursuing AI projects on their own. American and British conferees thought perhaps the Japanese had tried to do too much too soon.³⁰ The key to applying AI technology effectively is to establish realistic goals based on understanding what it can and can't do; what it is and what it is not. Hopefully, this handbook has provided a helpful beginning.

NOTES

CHAPTER 3

1. Robert E. Kahn "A New Generation in Computing," Spectrum, November 1983, 36.
2. Edward Feigenbaum and Pamela McCorduck, The Fifth Generation (Reading, Mass.: Addison-Wesley Pub., Co., 1983), 14.
3. Ibid., 18.
4. Kahn, "A New Generation in Computing," 38.
5. Trudy E. Bell, "Strategies," Spectrum, November 1983, 45.
6. Feigenbaum, Fifth Generation, 2, 13.
7. Tohru Moto-oka, "Japan," Spectrum, November 1983, 48.
8. Ibid., 50.
9. Feigenbaum, Fifth Generation, 126.
10. Mark A. Fischetti, "The United States," Spectrum, November 1983, 51.
11. Robert S. Cooper and Robert E. Kahn, "SCS: Toward Supersmart Computers for the Military," Spectrum, November 1983, 54.
12. Fischetti, "The United States," 53.
13. Microelectronics and Computer Technology Corporation, Information Brochure, 1984.
14. Brian W. Oaklev, "Great Britain," Spectrum, November 1983, 70.
15. Horst Nasko, "European Common Market," Spectrum, November 1983.
16. Paul Walton and Paul Tate, "Soviets Aim for 5th Generation," Datamation, 1 July 1984, 53-64.
17. Feigenbaum, Fifth Generation, 134.

18. Business Week, "AI the Second Computer Age Begins," 8 March 1982, 75.
19. James B. Schultz, "Weapons That Think," Defense Electronics, January 1983, 79.
20. Thomas D. Garvey and John D. Lowrance, "An AI Approach to Information Fusion," Journal of Electronic Defense, July 1984, 31.
21. J. P. Woodward and E. J. Cupples, "Selected Military Applications of Automatic Speech Recognition Technology," IEEE Communications Magazine, December 1983, 38.
22. Alan L. Meyrowitz, "Artificial Intelligence Research: Directions," Computers and People, March-April 1983, 14, 15.
23. Science, "The Pentagon's Ambitious Computer Plan," Vol. 222, 16 December 1983, 1213.
24. Business Week, "Artificial Intelligence is Here," 9 July 1984, 55.
25. Charles A. Pratt, "An Artificially Intelligent Locomotive Mechanic," Simulation, January 1984, 40.
26. Geoffry von Limbach and Michael B. Taylor, "Expert System Rules Read Natural Language," Systems and Software, August 1984, 119.
27. Patricia Keefe, "Natural Language Program Gaining Attention," Computer World, 28 November 1983, 77.
28. Tom Murphv "Artificial Intelligence Topics at IBM," Simulation, January 1985, 33.
29. Newsweek, "The Race to Build a Supercomputer," 4 July 1983, 58.
30. Austin American, "Closing the Computer Gap," 17 December 1984, E1.

ATTACHMENT 1

BIBLIOGRAPHY

A wealth of available literature discusses all aspects of artificial intelligence. Most of it, however, is written by computer specialists for other technical personnel. The following recommended bibliography is intended to provide a nontechnical introduction to an admittedly complex subject. While some of these volumes contain technical chapters, they also provide good introductory sections.

Dreyfus, Hubert. What Computers Can't Do, New York: Harper and Row Publishers, 1979.

Feigenbaum, Edward A., and McCorduck, Pamela. The Fifth Generation, Reading, Mass.: Addison-Wesley Publishing Co., 1983.

Gevarter, William B. An Overview of Artificial Intelligence and Robotics, Vol. I-Artificial Intelligence, Parts A and B, Washington, D.C.: National Aeronautics and Space Administration, 1983.

McCorduck, Pamela. Machines Who Think, San Francisco: W. H. Freeman and Co., 1979.

Spectrum, Vol. 20, November 1983. "Tomorrow's Computers" (Special Issue on Artificial Intelligence).

Walker, Roger S. Understanding Computer Science, Dallas: Texas Instruments Inc., 1981.

Winston, Patrick H. Artificial Intelligence, Reading, Mass.: Addison-Wesley Publishing Co., 1977.